

Cyberinfrastructure for Engineering June 5-6, 2003, Arlington, Virginia

This document summarizes the results of the Cyberinfrastructure for Engineering Research and Education workshop held by the National Science Foundation in June, 2003 in Arlington, Virginia.

The purpose of this workshop was to provide input from the engineering community to the Directorate of Engineering on how to make the best investments that support the creation and use of new information technologies in support of engineering research and education over the next decade.

Table of Contents

WORKSHOP ATTENDEES	2
TECHNICAL BACKGROUND	3
CYBERINFRASTRUCTURE AS INFRASTRUCTURE	3
THE NATURE OF CYBERINFRASTRUCTURE	4
THE VALUE OF SHARING OUR DATA	4
APPLICATIONS DEVELOPMENT: MIDDLEWARE OR MUDDLEWARE?	6
THE HARDENING AND MONITORING OF CYBERINFRASTRUCTURE	7
EDUCATION, OUTREACH, AND TRAINING	8
NSF ENGINEERING CYBERINFRASTRUCTURE CASE STUDIES	9
NEES	9
CLEANER	10
NNIN	10
SPECIFIC WORKSHOP TOPICS	11
CYBERINFRASTRUCTURE INVESTMENTS	11
CYBERINFRASTRUCTURE RESEARCH AND EDUCATION ACTIVITIES	12
CYBERINFRASTRUCTURE PARTNERSHIPS	12
CYBERINFRASTRUCTURE INVESTMENT ROADMAP	13
CYBERINFRASTRUCTURE ORGANIZATIONAL STRUCTURES	14
CYBERINFRASTRUCTURE PATH FORWARD	15
REFERENCES	16

Workshop Attendees

Richard Alkire	<i>University of Illinois, Urbana-Champaign</i>
Massoud Amin	<i>U. Minnesota Twin Cities</i>
Yigal Arens	<i>University of Southern California.</i>
Manohar Asnani	<i>Satyam Computer Services</i>
Ted Belytschko	<i>Northwestern University</i>
James E. Bernard	<i>Iowa State University</i>
David Bernstein	<i>James Madison University</i>
Patrick L. Brezonik	<i>University of Minnesota</i>
Randy Butler	<i>UIUC-NCSA</i>
Shiyi Chen	<i>Johns Hopkins University</i>
Dominic Di Toro	<i>University of Delaware.</i>
John Doyle	<i>California Institute of Technology.</i>
Kelvin K. Droegemeier	<i>University of Oklahoma.</i>
Gregory L. Fennes	<i>University of California, Berkeley</i>
Jose A. B. Fortes	<i>University of Florida.</i>
James H. Garrett, Jr	<i>Carnegie Mellon University</i>
Thomas R. Harris	<i>Vanderbilt University</i>
Larry P. Howard	<i>Vanderbilt University</i>
Marianthi Ierapetritou	<i>Rutgers University</i>
Thomas J. Impelluso	<i>San Diego State University</i>
Douglas W. Jacobson	<i>Iowa State University</i>
Grace Y. Lin	<i>IBM</i>
Mark Lundstrom	<i>Purdue University</i>
Michael Malone	<i>University of Massachusetts Amherst</i>
Barbara Minsker	<i>University of Illinois Urbana-Champaign</i>
Kyran D. Mish	<i>University of Oklahoma</i>
C. L. Max Nikias	<i>University of Southern California</i>
Cherri M. Pancake	<i>Oregon State University</i>
Ellen M. Rathje	<i>The University of Texas at Austin</i>
Andrei Reinhorn	<i>The University at Buffalo</i>
Don Riley	<i>University of Maryland</i>
Arthur C. Sanderson	<i>Rensselaer Polytechnic University</i>
Karsten Schwan	<i>Georgia Institute of Technology</i>
Frieder Seible	<i>University of California San Diego</i>
Jennie Si	<i>Arizona State University</i>
Janos Sztipanovits	<i>Vanderbilt University</i>
Sandip Tiwari	<i>Cornell University</i>


Technical Background

The content below represents a synthesis of presentation materials and technical background documents presented during the workshop. This content provides the context for the specific recommendations developed during the workshop breakout sessions, which are presented later in this workshop summary document.


Cyberinfrastructure as Infrastructure

Infrastructure is a fundamentally important part of what makes the United States the leading nation in the world. The engineered systems that power our economy and support our lives represent the unique confluence of a national dedication to developing a high-quality shared infrastructure, the availability of key natural resources, a skilled national work force, and the excellence of engineering practice within this country.

The National Science Foundation is now poised to begin the funding, design, development, and deployment of a national cyberinfrastructure that will provide the information-technology equivalent of our existing national infrastructural systems. It is imperative that this new cyberinfrastructure exhibit the same peerless and pervasive qualities as our existing national infrastructure, so that the United States can continue to lead the world in both the physical and the virtual realms.



Infrastructure: San Francisco-Oakland Bay Bridge
(image courtesy of LLNL)



Cyberinfrastructure: ASCI Red Supercomputer (image
courtesy of Sandia National Laboratory)

Infrastructure is a product of the engineering profession. The very qualities that define the nature of infrastructure (e.g., that it is ubiquitous, reliable, usable, etc.) are in large part due to the engineering sensibilities that underlie its design and deployment. So just as the term “cyberinfrastructure” has meaning in terms of its relation to “infrastructure”, the successful development of a national cyberinfrastructure will hinge on its design and deployment in relation to how the engineering community has insured the effective design and deployment of our shared national infrastructure.

The Nature of Cyberinfrastructure

Cyberinfrastructure is as ineffable in its description as it is essential in its utility. The *Report of the NSF Blue-Ribbon Panel on Cyberinfrastructure* (Atkins, et al, 2003) includes an entire appendix titled “More About What is Cyberinfrastructure”, which constitutes an implicit admission of the difficulty of precisely (or concisely) defining this key technological goal. The long-term requirements for national cyberinfrastructure are in large part unknown *a priori*, so they cannot be estimated with complete accuracy without concomitant incremental efforts at construction of this electronic infrastructure. Where much of our current infrastructure is compartmentalized among various engineering disciplines, cyberinfrastructure offers the promise of general abstract solutions that can be utilized (and constructed) by the entire engineering profession. Such general technologies offer the promise of unifying diverse professional practices in engineering as they improve the overall quality of engineering research and development. Cyberinfrastructure includes the following information technology components:

- Pervasive networks that link wired and wireless communications devices
- Centralized facilities for high-performance computing, shared data access, and generation of experimental results
- Software components that aid in domain-specific applications development, human-computer interaction, data replication and transfer, and remote authentication
- New hardware and software technologies to support synthesis and fusion of data from various experimental and computational sources
- Novel technologies (including embedded systems) to permit pervasive access to computing and data resources
- Accurate and reliable sensor technologies, including self-configuring networks of intelligent sensors and sensor systems for rapid estimation of mechanical, electrical, chemical, and biological responses

While this is not an exhaustive list of the capabilities for new cyberinfrastructure, these components are among those that will have the most impact on the engineering profession in particular, and on the quality of life in our nation in general.

The Value of Sharing Our Data

Errors using inadequate data are much less than those using no data at all

- Charles Babbage -

The development of a national cyberinfrastructure constitutes a third stage in the National Science Foundation’s investment in information and computing technologies. The first stage was oriented towards centralized supercomputer facilities, as NSF’s investments were primarily oriented towards “big iron” computing, a strategy that led to the creation of the various NSF supercomputer centers. The second stage is the current state of investment, with a breadth of novel applications funded via NSF/CISE, large new centers funded by MREFC resources, and many other computing and information efforts

distributed across the breadth of the Foundation's technology investment portfolio. This stage is characterized by investments in advanced networking technologies (including middleware that abstracts network operations such as transport and authentication), and hence it might be referred to as the "big pipe" stage of investment.

The future of cyberinfrastructure is increasingly seen in a different direction, one defined by the sharing and discovery of data, including large amounts of distributed information. Transparent access to any amount of relevant data (e.g., access which can be gleaned currently using existing web searching and browsing techniques) has improved the productive abilities of our scientists and engineers -- one can only imagine what new discoveries await our ability to utilize our cyberinfrastructure to discover those needles of wisdom that lie in the haystacks of data that will be available once our scientific and engineering communities can fully embrace the notion of sharing data and information. This third *data-centric* ("big data") stage of NSF investment is what will finally permit our national information technology investments to become a shared infrastructure instead of a separable set of research and development components.

The engineering profession relies much more on its community-generated data than any other profession, because engineers learn from experience, past mistakes, theories, simulations and models (as well as the fusion of these various forms of collected wisdom). Cyberinfrastructure will aid engineers in collecting, validating, transmitting, storing, searching, displaying, analyzing, and archiving more and better data in increasingly faster and more reliable ways, so as to shorten the time from data collection to information and from information to knowledge that improves engineering.

Only if the flow and management of engineering data becomes so transparent that it becomes a commonplace practice which all engineers may then take for granted, can we refer to a successful cyberinfrastructure for Engineering. While this ultimate end may not directly result in *a priori* quantifiable metrics, it is an ambitious and productive goal for the greater engineering community to strive for. Quantifiable metrics can be established *a posteriori* based on usage and access frequencies.

It is essential that NSF/Engineering's investments in cyberinfrastructure must implicitly recognize the unique and productive nature of engineering culture, and how this professional culture differs from that of the physical and biological sciences. Where science is generally motivated by the desire to *discover* new information about the physical world, engineering is built on the *integration* of information into novel applications of technology that improve the quality of the human experience. Thus science is a culture of discovery whose outcome is scientific truth, where engineering is a professional culture of integration whose outcome is technology. The most productive applications of cyberinfrastructure to engineering practice will require the fusion and accumulation of data in a manner that may be unique to the engineering profession.

Finally, future cyberinfrastructure resources should be accessed and used interactively (not in batch mode, nor reserved), thereby prompting "big-data" research in resource or service trading and giving all end users the virtual proximity to resources they need in order to have apparent individual ownership of and direct access to their information and data. The idea here is to prevent the 'big iron remote supercomputer center' syndrome that has caused many research groups to build their own facilities, and also to encourage

groups to share resources. The result should be that end users gain significant improvements in their ability to conduct cutting-edge research.

The implication on technical cyberinfrastructure research is that substantial improvements are required in the ability for distributed systems to deliver high levels of performance and of reliable and trusted service. This means that it is important to support hardware and systems research that better enables systems to interact, to be configured to better interact on specific tasks, to be self-healing, to better match current needs or conditions. *Only when such robust and reliable information capabilities are present will the resulting technology deserve the characterization as infrastructure.*

Applications Development: Middleware or Muddleware?

Software will be an essential component of any successful national cyberinfrastructure. In this venue of software development the expertise of the newest engineering discipline, *software engineering*, must be brought to bear towards the goal of cost-effective design and deployment of cyberinfrastructure. The software engineering profession has its roots in past attempts at creating a shared information infrastructure for research collaboration (e.g., the Multics debacle), and the lessons learned from past failures of large-scale software development must be applied in order to insure future successes.

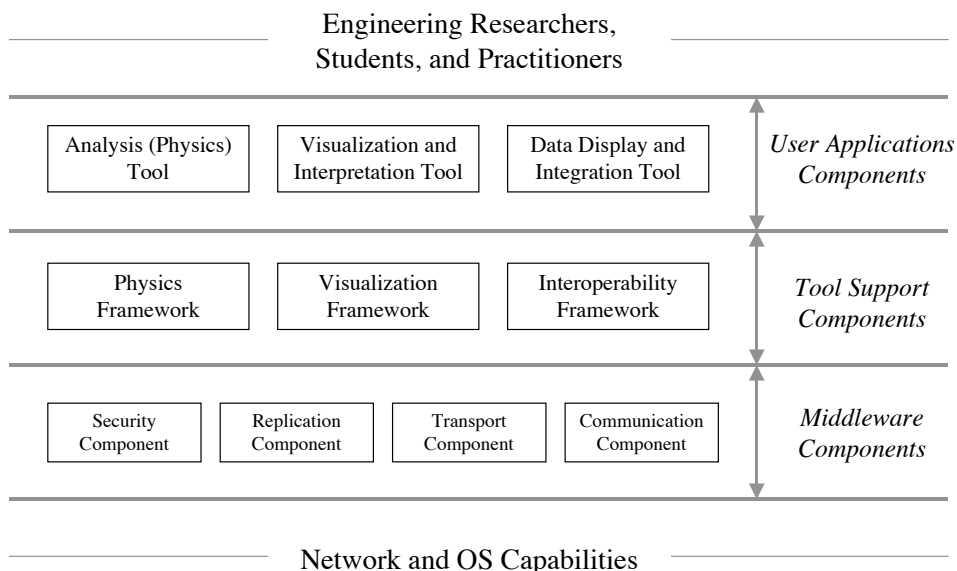
The National Science Foundation currently funds a substantial amount of software development in support of cyberinfrastructure, under specific programs such as the National Middleware Initiative (NMI) or using targeted information-technology grants such as Information Technology Research (ITR). The disparate research efforts are fundamentally important to the successful deployment of a national cyberinfrastructure, but their utility is often compromised by their exploratory nature, as research efforts in infrastructure are seldom appropriate for immediate deployment as production systems relevant to an infrastructural capability. The Foundation needs to develop (either internally or in coordination with federal mission-oriented agencies) appropriate reduction-to-practice funding and coordination mechanisms so that the most promising exploratory software technologies can be tuned, hardened and deployed in a manner likely to make our national cyberinfrastructure as ubiquitous and reliable as our shared physical infrastructure. Past experience has shown that the failure to coordinate, tune and harden research software efforts will result in a hopeless muddle of non-interoperable software that can never be considered as a reliable information infrastructure.

The software that is utilized to develop engineering applications for cyberinfrastructure provides a promise of unifying the methods and practices of the engineering profession. Currently, engineering practices are Balkanized, with each engineering discipline basing its professional practices on oft-wildly different metrics: for example, the code-based practices in the Civil Engineering profession are largely oriented towards a standards-based approach to mitigate the potential risk of litigation, where the practices used in commercial electrical engineering profession (e.g., microprocessor design) are closer to those of emerging scientific fields, e.g., best-practices based on the newest technology.

The agile, generic and mutable characteristics of software and information technology provides a rare opportunity to unify many of the diverse practices found in the engineering profession. For example, computational analysis techniques appropriate for

antenna design are strikingly similar to those used to analyze long-span bridges or high-performance aircraft -- it is thus through the development of enabling *general-purpose engineering applications frameworks* that cyberinfrastructure investments hold the promise of a more unified and systems-oriented approach to the whole of engineering practice.

The architecture of such frameworks is shown in the figure below, where these applications frameworks can be seen as mediating the interface between lower-level computer science technology and higher-level applications needs. Since this interface is susceptible to the most rapid pace of technological change (e.g., as user requirements emerge over time, and as engineering practice matures), it is essential to develop software components that insulate the engineer from the technological details of middleware and lower layers. *Insulating the long-term functional needs of the community from the underlying technological implementation details is an essential aspect of physical infrastructure, and hence this insulating role must be seen as an equally important component of any cyberinfrastructure plan.*



In the figure above, applications frameworks are represented for engineering analysis, for visualization, and for interoperability of common software tools. These are not intended to form an exhaustive list of needs, but instead to represent the most obvious instances of common unifying engineering applications.

The Hardening and Monitoring of Cyberinfrastructure

The United States is the most advanced and powerful nation on earth in large part because of the quality and reliability of our shared infrastructure. This makes our national infrastructure both a unique asset to our citizens, and an inviting target to our enemies. The events of September 11, 2001 have shown the need for hardening our physical infrastructure, and we need to insure that any new national cyberinfrastructure is

equally secure, lest we invite our enemies to inflict serious damage upon our country by destroying our capability to utilize our considerable information technology resources.

Hardening and protection of our new cyberinfrastructure will be a difficult undertaking, because many of the means available to protect infrastructure assets will not necessarily insure the safety of analogous information resources. For example, we currently restrict access to critical structures such as airport terminals by means of proximity constraints, yet the entire notion of proximity has less relevance in cyberspace, where hackers and information warriors can readily span national and international boundaries while attempting to wreak damage on our national cyberinfrastructural resources.

It is clear that investments in development of cyberinfrastructure will require concomitant investments in hardening and monitoring of the health of that information infrastructure, so that our national cyberinfrastructure can remain reliable and resilient in the face of natural and man-made hazards. In particular, hardening efforts that will preclude the form of asymmetric warfare that has characterized the nation's war on terrorism must be given high priorities, especially when those new cyberinfrastructural components will support society-critical functions such as transportation, communications, power, water, and financial systems.

Education, Outreach, and Training

The construction of a national cyberinfrastructure requires a technologically-skilled workforce, and it also provides a unique opportunity for training this workforce using the best practices of both learning science and information technology. In order to design and deploy a national cyberinfrastructure system, a cadre of technically-savvy personnel must be deployed to work in collaborative fashion within (and at the interfaces of) many diverse fields, including:

- Systems engineering
- Information science and engineering
- Computer science, computational science, and computational engineering
- Software engineering and software project management
- Learning science and educational practice

This cadre of technology experts must include operational personnel (e.g., software and tool developers and maintainers, both for generic support software and for domain-specific tools). Furthermore, these professionals must be linked to key cyberinfrastructure research and usage efforts, so that intellectual advances and insights into future needs are continuously infused into the cyberinfrastructure initiative. The implication here for cyberinfrastructure research is that we need to better understand how to represent new services, how to adapt them to user needs and platform capabilities, how to deploy such services into different levels of the cyberinfrastructure (network, operating system, middleware, applications), and how to integrate services that operate at different levels of abstraction to ensure desired levels of quality (e.g., performance, trust, reliability).

Recent advances in the learning science have proven to be of import for the education of engineers (Harris, et al, 2002, Bransford, et al, 1999). Research has shown that creating learning environments that are more learner centered, knowledge centered, assessment centered and community centered can improve learning in engineering courses and can move such education toward the creation of engineers who are adaptive experts. Learning technologies supported by well-designed cyberinfrastructure can improve the efficacy of such education and significantly improve its efficiency. Key concepts that improve engineering education are a movement toward “challenge-based” instruction in which engineering subject teaching is placed in context through a series of appealing challenges, and substantial increases in the efficiency of formative assessment. Technologies for developing courseware based on these ideas, construction of repositories for digital materials and systems for course management are examples of key technologies that enable educational reform.

NSF Engineering Cyberinfrastructure Case Studies

NSF/Engineering is already associated with several key cyberinfrastructure investments, including efforts in earthquake engineering, environmental engineering, and nanotechnology. This technology portfolio is summarized below.

NEES

NEES is the acronym for the George E. Brown, Jr. Network for Earthquake Engineering Simulation Major Research Expenditure (MRE). This first-ever NSF/Engineering MRE is intended to provide a geographically distributed earthquake engineering experimental and computational resource system that will enable novel forms of collaborative research in earthquake engineering. The NEES project consists of fifteen university experimental awards in structural, geotechnical, and tsunami engineering, a systems integration award to provide the cyberinfrastructure that binds the experimental systems together with a computational subsystem, and an award to develop the NEES Consortium that will accept and administer the shared-use system for the period Oct 1, 2004 through Sept 31, 2014.

The NEES project represents an excellent case study in the deployment of cyberinfrastructure and the lessons to be learned from early experiences. The key lesson is that cyberinfrastructure is fundamentally a "human problem." The most difficult challenges stem from the diversity of users, their broad range of skill sets and familiarity with advanced IT, the broad spectrum of needs to be addressed, the fact that most participants have had little experience sharing data or facilities, and the extremely high expectations of all parties concerned. A second lesson is that computer scientists can easily misunderstand or underestimate the priorities of the user community, which can lead to wasted effort and tension between the two groups. For example, the NEES IT specialists initially assumed that their focus should be on development of new IT capabilities (telepresence, collaboration technology, security). The earthquake engineers, however, had other priorities (chiefly data accessibility, usability, and system reliability). This meant that time was spent developing technology features that were low priority for users, while the most important targets were neglected. Course corrections were made, but the misunderstandings affected both the schedule and the spirit of cooperation between earthquake engineers and computer scientists. Since cyberinfrastructure requires

the participation of both, it is important to find ways to avoid the friction that occurs when technology is deployed without sufficient attention to the requirements of the target user community.

CLEANER

Environmental cyberinfrastructure is critical to our national goals and aspirations. The development of fundamental scientific knowledge coupled to the understanding and management of engineered systems will be essential to shape policy and assess implications of key decisions.

The Engineering Directorate at NSF is developing the **CLEANER** (Collaborative Large-scale **E**ngineering Assessment Network for **E**nvironmental **R**esearch) initiative to address such issues. Three community-based workshops have defined research themes across a set of common problem domains, such as river and coastal systems, and urban airsheds, across a region impacted by complex, interacting phenomena. These large-scale problems are of fundamental importance to human health, economic development, national defense and other areas where environmental understanding and management play a pivotal role in our ability to observe, predict, and respond to vital issues. The information science and technology to enable such monitoring and prediction is not currently available. These challenges define fundamental research issues in engineering cyberinfrastructure including the availability of distributed, real-time, embedded information systems, as well as large scale modeling, simulation, and data management.

In addition, such problems entail the need for engineering analysis, development of risk assessment, and adaptive management. Such large-scale problems cannot be addressed by conventional individual research projects.

As the workshops have defined this initiative, **CLEANER** will be *a networked cyberinfrastructure of environmental field facilities that enables formulation and development of engineering and policy options for the restoration and protection of environmental resources*. The goal of **CLEANER** is to provide options for maintaining and improving the environment through understanding and predicting the behavior of anthropogenically-stressed regional environmental systems (RES), utilizing large-scale projects involving several environmental field facilities (EFF) collaborating with related large-scale projects in an “Engineering Assessment Network” (*EAN*).

NNIN

The National Nanotechnology Infrastructure Network (NNIN) as an integrated national network of user facilities that will support the future infrastructure needs for research and education in the burgeoning nanoscale science and engineering field. The facilities comprising this network will be diverse both in capabilities and research areas served as well as in geographic locations, and the network will have the flexibility to grow or reconfigure as needs arise. The NNIN will broadly support nanotechnology activities outlined in the National Nanotechnology Initiative investment strategy.

Specific Workshop Topics

The content below presents summaries of the discussions from the various breakout groups at the workshop. The reader is invited to supplement the content below with the summary discussion content on the NSF web pages, and to review the various individual Personal Points of View (PPV) that were submitted by workshop attendees.

Cyberinfrastructure Investments

Future cyberinfrastructure investments should proceed from the premise that *evolutionary* improvements to current information systems will not be sufficient to promote an appropriately robust and reliable cyberinfrastructure: instead, *revolutionary* improvements are required to support the “big-data” model required for future engineering integration and scientific discovery processes. It is thus essential to infuse the cyberinfrastructure technology investment portfolio with systems engineering sensibilities, so that academic computer science research results can be supplemented with appropriate control and systems theories to insure that the resulting information infrastructure will be reconfigurable, reliable and robust enough for future applications.

Furthermore, given the proliferation of distributed devices utilizing embedded technologies, it is essential that national cyberinfrastructure investments must include substantial commitments to the analysis, design, and deployment of distributed systems that utilize wireless communications, embedded processors, and realtime processing. These pervasive technology investments should serve as a counterweight to existing NSF investments in centralized high-performance computing centers, and will serve to diversify NSF’s technology investment portfolio in cyberinfrastructure venues.

In the context of the cyberinfrastructure, engineering should not merely be an end user that provides only domain specific applications and specifies domain requirements for cyberinfrastructure research and development. Instead, engineers should play an integral role in partnership with key infrastructure providers and application end users, so as to address cyberinfrastructure systematically. The need for cyber-infrastructure is not just that of procuring more hardware: it instead embraces many important aspects of engineering approaches in dealing with congestion, including control, scheduling, stability, security, service, performance, and many other important features of a seamless cyberinfrastructure environment.

Investments must be made that recognize the unique characteristics of the proposed new cyberinfrastructure, e.g., “big-data” applications will be facilitated by fundamental research in human-computer interaction (including visualization as well as other senses beyond that of sight), and in development of agile software tools that permit rapid development of novel data-handling capabilities beyond current mining and filtering operations. New applications that fully realize the capability of humans to understand rich sources of information will be a necessary driver for a well-developed national cyberinfrastructure.

Cyberinfrastructure Research and Education Activities

All areas/disciplines of engineering will benefit from the cyberinfrastructure in the form of new discoveries and advances in research and education. Advances will come from data searching and sharing (including legacy data), remote participation, large scale simulation, and visualization. The required characteristics are that data access and management in the broadest sense, including simulation and visualization need to be pervasive, fast, reliable, and usable. It needs to develop into a commonplace activity such as other infrastructure services (e.g., checking e-mail) have already become.

Major challenges are to have data (with domain specific attributes) archived and stored that it can be searched by information queries. The establishment and curation of these domain specific data sets are paramount for the success of an engineering cyberinfrastructure. It is not clear where the sustained support for this critical function will come from and if this can be accomplished in a decentralized fashion.

Currently we have Computer Scientists and Domain Engineers, and what is needed is a facilitation capability to bring these two groups together. This is an important function that is critical to the success of an Engineering cyberinfrastructure. It is not clear if this facilitation can happen in the virtual space or if physical/personal facilitation is required. Certainly experience from NEES shows that it works much better as soon as personal facilitation is involved. Once the engineering cyberinfrastructure is design, it will require a special type of professional to implement and maintain such an infrastructure. This professional needs to be proficient in both the IT infrastructure and the domain needs.

Cyberinfrastructure Partnerships

Engineers need to articulate the challenges inherent in the development of a national cyberinfrastructure to make these challenges specific to engineering peers, but also in a way that can be understood by computer scientists: this will permit both engineers and computer scientist to seek and find appropriate enabling technologies. Better abstraction of engineering problems for communications with computer scientists is critical and characteristic of cyberinfrastructure. For example, rapid prototyping tools and visualization, commonly are used in computer science community, are examples for engineers to learn to communicate with a larger community outside our engineering fields. Engineers need to work with computer scientists to develop computer tools across disciplines with increased ease or adaptation. This collaboration would also reduce and prevent redundant work, enable more sharing of results.

Engineers and computer scientists have complementary skills. Cyberinfrastructure enhances intellectual diversity, working with other communities. It takes time, tools, capabilities, and discipline to develop strong collaboration. However, interdisciplinary approaches are keys to the effective deployment of a shared cyberinfrastructure (just as they have been in the effective deployment of our national physical infrastructure), as long as the underlying partnership is led from the applications side, as this is where the user requirements are to be found. It takes a strong partnership of several communities, with engineering and computer science playing important roles, to address research problems in multiscale chemistry, composite materials, turbulence, civil infrastructure,

driver behavior, air quality, etc. cyberinfrastructure is an opportunity for researchers to work together to address very large problems that have not been able to address without heterogeneous computing, experimental, analytical facilities, all in an integrated manner.

There need to be significant incentives for both engineering and computer science parties to come together in order to pursue a common goal. These incentives can be in the form of funding opportunities and/or special academic recognition of cross-disciplinary activities. In addition, these partnerships need to be extended to include industry in order to be successful in the development of cyberinfrastructure. A large amount of engineering data is in and comes from industry and regulatory organizations and requires just as much management as research data. If industry identifies the need for a new breed of professional, namely the domain IT specialist, the partnership at the university level will also be driven by new education programs.

The cyberinfrastructure needs to be service oriented by catering to the user community. Unless the end-user is directly involved in the cyberinfrastructure development, it will never achieve the user friendliness and buy-in needed for broad outreach. This end-user orientation must be facilitated for effective partnerships to form in support of development of cyberinfrastructure.

Past successes in developing shared cyberinfrastructure (e.g., the NSF supercomputer centers) must be augmented by future partnerships that will include the emerging field of pervasive technology. An example of such a new partnership would be one that recognizes the importance of networks of wireless communications devices (e.g., our nation's cellular telephone infrastructure) that are dynamically reconfigurable, utilize embedded technologies, and whose requirements are incompletely known, since this is an emerging field of engineering technology.

Examples such as this one demonstrate that "new partnerships" can have two disparate meanings, namely new partnerships of existing technology stakeholders must be formed using past models of success, and partnerships leveraging emerging new technologies (such as those required for distributed realtime systems) must be formed as well.

Cyberinfrastructure Investment Roadmap

Technology exists only to serve the human condition. All technology development efforts must therefore begin with the user community for appropriate end-user initial requirements-gathering efforts. The history of information technology amply demonstrates that user requirements will evolve as new technology is deployed, because novel applications generate changes to original requirements: but regardless of this evolution in what the user communities may find desirable, successful evolution and tracking of user requirements always begins with an initial assessment of available user requirements, e.g., reliability, usability, etc. Furthermore, any and all technologies that aspire to be considered sufficiently ubiquitous and useful to warrant being termed *infrastructure* must consider first and foremost the needs of serving humanity.

Cyberinfrastructure planning must thus start with the user communities, and it must recognize the unique characteristics of modern professional practice, e.g., distributed teams of experts working collaboratively across geographical and disciplinary

boundaries. One of the most important early goals of cyberinfrastructure planning must therefore be to facilitate collaboration using information technology to bridge the distance between members of interdisciplinary teams.

Because information systems can be agile and mutable, it is worthwhile to seek commonalities in end-user requirements within the various user communities, so that the resulting information systems can be redeployed in a myriad of disparate settings. This generic capability of software provides two obvious benefits, namely:

- It provides cost-effective use of resources by exploiting economies of scale and the low marginal cost inherent in reusable software, and
- It helps user communities to recognize common modalities of practice, which permits them to become less fragmented, more inclusive, and more oriented towards the interdisciplinary systems-oriented problem-solving approaches that are required for effective practice in the future of science and engineering.

These commonalities need to be evaluated and the resulting assessments propagated back into the design and deployment practices used for developing a national cyberinfrastructure. This feedback between user and information technology communities will aid the latter in cost-effectively serving the real needs of the former.

Because our nation already possesses many elements of a cyberinfrastructure (e.g., large computer centers, wired and wireless communications capabilities), it is essential to begin the process of development of an investment roadmap for cyberinfrastructure with a frank assessment of available existing resources, so that the deployment of a national information capability can proceed incrementally by revision and modification of existing resources. In addition to an assessment of existing and required information components, specific resource allocations must be provided for the all-important tasks of integrating these components into useful and reliable systems. As the number of cyberinfrastructural components increases with time, the effort (and hence cost) of integrating among them can be expected to exhibit a concomitant growth, and appropriate financial and personnel resources for performing these integration efforts must be allocated in order for any national cyberinfrastructure deployment efforts to be achieve long-term success.

Cyberinfrastructure Organizational Structures

In the deployment of an agile, reliable, and extensible cyberinfrastructure, it is extremely unlikely that any organizational structure will constitute a “one size fits all” construct suitable for general use. There will be a need for a talented and reliable common personnel infrastructure of both specific domain and general technology specialists (e.g., a centralized source of consulting assistance for cyberinfrastructure development teams), but these repositories of information technology experts will need to address both the common service aspect of cyberinfrastructure and the details of domain-specific integration. The particulars of determining the best organizational constructs to provide these disparate services will need to be discovered as cyberinfrastructure is designed, developed, and deployed.

The past and present development of computational infrastructure (e.g., PACI, Terascale, etc.) has demonstrated the utility of centralized hubs of expertise where common core

infrastructure capabilities (e.g., HPC applications tuning, data modeling, etc.) can be concentrated for national use. Distributed structures for domain-specific support can take on many forms (including the central model used in current computational infrastructure), but these distributed teams should not be so fragmented that organizational control and overlap becomes problematic. The NEES project provides a case study of how a distributed model of project management may not work well in settings of academic partners, as there are seventeen awards that are administered by NSF, but none of them contain any overall responsibility for the project's systemic health.

Probably the most important goal of cyberinfrastructure organizational structures is that they will need to facilitate agile response to changing technological and social conditions, so that they will require flexibility and adaptability in their organizational structures. Technical organizations that have successfully managed long-term change (e.g., Hewlett-Packard, IBM, etc.) should be studied to learn how to deploy adaptable organizational constructs to support development of a national cyberinfrastructure.

Cyberinfrastructure Path Forward

The path forward for NSF/Engineering in developing a cyberinfrastructure plan includes these essential steps:

1. Identify all critical engineering domains and establish end user requirements for each domain. These critical areas of engineering interest should be identified at a sufficiently high level to permit enumeration of a manageable number of engineering domains, but they should also be specific enough to be able to identify key differences in underlying user requirements.
2. Identify and characterize potential commonalities among the individual engineering domain requirements so that common/core infrastructure strategies can be developed (i.e., those core areas where commonalities among the disparate engineering communities can be identified and exploited).
3. Identify existing NSF (and other federal) information technology resources that can be deployed to help provide support for those core strategies identified as candidate common technology areas in step (2).
4. Work within the current NSF/federal information technology infrastructure to provide the required core facilities and the flexibility to support domain specific investments in engineering information technology, using those existing resources where appropriate, and with new investments developed as required to fill the gaps in the existing technology portfolio or to support novel technologies that are likely to be capable of supporting other core engineering cyberinfrastructure needs in the future.
5. Do not limit the implementation to communities of research and education only: the larger community of practicing engineers is an invaluable resources to be tapped for development of a new national cyberinfrastructure, and the engineering community is unique in this connection to a vast body of relevant expertise in design, development, deployment, tuning, hardening, and financing of the required new technologies.

6. Develop a feedback strategy to facilitate the integration of core infrastructure efforts (e.g., middleware) with domain infrastructure (e.g., applications frameworks), so that the lessons learned dealing with users in the domain-specific venues can be propagated back into the core infrastructure efforts, so that our national cyberinfrastructure can be deployed in an appropriately agile manner.

References

- Atkins, D.E, et al, (2003), “Revolutionizing Science and Engineering through Cyberinfrastructure: Report of the National Science Foundation Advisory Panel on Cyberinfrastructure”.
- Harris, T. R., Bransford, J. D. and Brophy, S. P. (2002). "Roles for learning sciences and learning technologies in biomedical engineering". *Annual Review of Biomedical Engineering*, vol. 4, pp. 29-48.
- Bransford, J.D., Brown, A.L. & Cocking, R.R. *How people learn: Brain, mind, experience, and school*. Washington, DC: National Academy